

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

EV369764933

**Identification of Relationships In An Environment**

Inventors:

Shelly D. Farnham

William L. Portnoy

Lili Cheng

ATTORNEY DOCKET NO. MS1-1914US

## **TECHNICAL FIELD**

The systems and methods described herein relate to identifying one or more relationships between two points in an environment, such as a social environment.

## **BACKGROUND**

Computer systems are continuing to grow in popularity and are frequently interconnected with other computer systems via networks, such as local area networks (LANs) and the Internet. Features such as electronic mail (email), instant messaging, and project collaboration encourage the use of computer systems coupled to networks. These features allow users to, for example, communicate with other users, retrieve information, and share common documents.

In some situations, a user may need to interact with an unfamiliar person, an unfamiliar department, or an unfamiliar group. In other situations, a user may desire to learn how another person, department, or group is related to the user. For example, a user may want to talk to another person in an organization, but has never been introduced to that other person. In this situation, the user would like to know if there is a common person or a common group with which both the user and the other person are associated. In another example, a user may want to learn about a particular project, but doesn't know if they have any relationship to the project.

Attempting to discover these types of relationships manually is time-consuming and inefficient. For example, if a user asks a large group of people whether they know a particular person, that user spends a great deal of their time

1 communicating with these people, and takes time away from each person that is  
2 contacted. Further, attempting to manually search through various organizational  
3 charts, mailing lists, and other information to discover a relationship between two  
4 people (or between a person and a group) is time-consuming and may not  
5 accurately discover all relationships.

6 It would be desirable to provide an improved approach to identifying one or  
7 more relationships between two points in an environment.

## 8 9 10 **SUMMARY**

11 The systems and methods described herein identify and display  
12 relationships between points in an environment. In a particular embodiment, a  
13 method identifies components associated with a first end point in an environment  
14 and identifies components associated with a second end point in an environment.  
15 A determination is made regarding whether any of the identified components are  
16 associated with both the first end point and the second end point. A display is  
17 generated that shows relationships between the first end point, the second end  
18 point, and any components associated with both the first end point and the second  
19 end point.  
20  
21  
22  
23  
24  
25

## **BRIEF DESCRIPTION OF THE DRAWINGS**

Similar reference numbers are used throughout the figures to reference like components and/or features.

Fig. 1 illustrates an example environment containing various social contexts and points within those contexts.

Fig. 2 is a flow diagram illustrating an embodiment of a procedure for identifying one or more relationships between two points in an environment.

Fig. 3 is a flow diagram illustrating an embodiment of a procedure for displaying one or more relationships associated with points in an environment.

Fig. 4 illustrates an example graphical display of information related to multiple points in a social environment.

Fig. 5 illustrates the example graphical display of Fig. 4 with additional information displayed regarding a particular point in the social environment.

Fig. 6 illustrates another example graphical display of information related to multiple points in a social environment.

Fig. 7 illustrates another example graphical display of information related to multiple points in a social environment.

Fig. 8 illustrates a general computer environment, which can be used to implement the techniques described herein.

## **DETAILED DESCRIPTION**

The systems and methods discussed herein identify and display one or more relationships between two points in an environment, such as a social environment. Social information is modeled according to a computer networking approach. Users are provided information in a graphical form representing, for example, relationships and patterns in their social spaces or social networks. The systems and methods discussed herein analyzes relationships associated with two end points (e.g., users) to identify one or more common relationships between the two end points. These common relationships may include common people or common groups that “connect” the two end points within an environment, such as a corporate environment. Relationships can be identified by analyzing organizational charts, group information, mailing lists, and the like.

As used herein, the term “component” may include a person, a project, a group of persons, a department, an entity, an association, and the like. A component may also be referred to as a “node” or a “point”. An “end point” can be any component, such as a person, a group, a project, or a department. As used herein, the terms “environment”, “computing environment”, “corporate environment”, “social environment”, and “social network” are used interchangeably. As used herein, a “mailing list” may also be referred to as a “distribution list”.

Fig. 1 illustrates an example environment 100 containing various social contexts and points within those contexts. Environment 100 may also be referred to as a “social environment” or a “social network”. A first social context 102 is labeled “Point A Social Context” and a second social context 106 is labeled “Point B Social Context”. Another social context 104 overlaps a portion of social

1 contexts 102 and 106, and is labeled “Overlapping Social Context”. A particular  
2 social context includes any number of components, such as users that are related to  
3 one another and groups or departments associated with those users. For example,  
4 a social context associated with a particular user may include other individuals  
5 with which the particular user is associated, such as friends, co-workers, other  
6 members of common groups or clubs, and other members of common mailing  
7 lists. In another example, a social context associated with a particular project may  
8 include other components associated with that project, such as employees assigned  
9 to the project, persons managing the project, related projects, and other entities or  
10 organizations related to the project. An example overlapping social context  
11 contains one or more components that are common to the two overlapped social  
12 contexts, such as a common person or a common project. A particular social  
13 context, including an overlapping social context, may contain any number of  
14 components.

15 In the example of Fig. 1, social context 102 includes “Point A” identified  
16 by reference number 108. Thus, social context 102 is the social context associated  
17 with Point A and labeled “Point A Social Context”. Additionally, social context  
18 106 includes “Point B” identified by reference number 110. Thus, social context  
19 106 is the social context associated with Point B and labeled “Point B Social  
20 Context”. Overlapping social context 104 includes one or more components that  
21 are common to both Point A Social Context 102 and Point B Social Context 106.  
22 Overlapping social context 104 may also include additional components that are  
23 not common to both Point A Social Context 102 and Point B Social Context 106.  
24 Additional details regarding social contexts are provided below.  
25

1        Fig. 2 is a flow diagram illustrating an embodiment of a procedure 200 for  
2 identifying one or more relationships between two points in an environment, such  
3 as a social environment. Initially, a first end point is identified (block 202) along  
4 with a second end point (block 204). For example, the first end point and the  
5 second end point may be identified by a user desiring to learn of relationships  
6 between the two end points. These end points may be users, groups of users,  
7 projects, or other components. The two end points may be different types of  
8 components, such as a user and a project, or a project and a department.

9        Procedure 200 continues by identifying components associated with the  
10 first end point (block 206). This identification can be performed, for example, by  
11 analyzing organizational charts to identify associated employees (e.g., managers or  
12 other employees in the same department), mailing lists to identify associated  
13 individuals, personal profiles, organization membership lists, project participant  
14 lists, club lists, group lists, email address books, and the like. For each identified  
15 component, the procedure determines strengths of links between the identified  
16 component and the first end point. The procedure also identifies components  
17 associated with the second end point (block 208). This identification can be  
18 performed, for example, by analyzing any of the items discussed above with  
19 respect to identifying components associated with the first end point. For each  
20 identified component, the procedure determines strengths of links between the  
21 identified component and the second end point.

22        The procedure then determines a path strength for each path between the  
23 two end points (block 210). A path between two end points may include any  
24 number of components and any number of links between those components. In a  
25 particular embodiment, any path that does not have a minimum path strength is

1 ignored or discarded. This minimum path strength threshold may be established  
2 by a network administrator or other user.

3 In a particular embodiment, the path strength for each path between the two  
4 end points is determined by first identifying the strength of each link between the  
5 two end points (e.g., the strength of the link between the first end point and a  
6 common component, and the strength of the link between the common component  
7 and the second end point). The strength of a particular link is rated between 0.0  
8 and 1.0, where 0.0 indicates no relationship between the two components (i.e., the  
9 two components at opposite ends of the link) and 1.0 indicates a strong  
10 relationship between the two components. Various procedures can be used to  
11 determine the strength of a particular link. For example, a link's strength may  
12 increase if the two components are identified on multiple group lists, mailing lists,  
13 etc. Additionally, a link's strength may increase if the two components are  
14 contained in one or more lists and those lists are relatively short (signifying a  
15 stronger relationship among members of the shorter list). In one embodiment,  
16 mailing lists, group lists, or other lists having more than a particular number of  
17 members (e.g., more than 100 members) are ignored for purposes of identifying  
18 relationships between two components. Another example of calculating the  
19 strength of links between components is described in U.S. Patent Application  
20 Publication No. US 2003/0167324 A1, published September 4, 2003, entitled  
21 "Social Mapping of Contacts from Computer Communication Information".

22 In certain embodiments, the strength of the links are normalized based on  
23 the number of mailing lists, group lists, or other lists in which each component is a  
24 member separately, such that the link strength is smaller if the components are  
25 each in a large number of lists.



1       The strengths of the links between the two end points are multiplied  
2 together to obtain a path strength. For example, if a particular path between two  
3 end points has two links, one with a link strength of 0.25 and the other with a link  
4 strength of 0.80, the resulting path strength is 0.20 (0.25 x 0.80).

5       Any path that contains at least one link with a strength of zero will have a  
6 path strength of zero, indicating that there is no common relationship between the  
7 two components at opposite ends of the path. Similarly, any link having a strength  
8 of zero indicates that there is no relationship between the components at opposite  
9 ends of the link.

10       The path strength provides an indicator of a relationship between two  
11 components. For example, if a common component is strongly associated with  
12 both a first end point and a second end point, that path will have a relatively high  
13 path strength. Examples of strong paths include 1) the situation where the  
14 common component is listed on the same mailing list as the two end points and the  
15 mailing list is relatively short, and 2) the situation where both end points are  
16 individuals that report to the same manager and the manager is the common  
17 component.

18       The procedure continues by ranking the paths between the first end point  
19 and the second end point (block 212). This ranking of paths may be based on the  
20 path strengths or other information. Finally, the relationships between the two end  
21 points are displayed graphically to a user (block 214), such as a user desiring to  
22 learn of relationships between the two end points. The relationships between the  
23 two end points may include one or more links (e.g., connecting lines) that  
24 represent relationships between the components at each end of the link. For  
25 example, a link is displayed between two people that work in the same department

1 or are listed in a common mailing list. Additionally, links may be established  
2 between a group and each member of that group that is included in the graphical  
3 display. Additional details regarding the display of the relationships are provided  
4 below with reference to Fig. 3. In certain embodiments, only the strongest paths  
5 between two endpoints are displayed, such as the three strongest paths or the five  
6 strongest paths. These embodiments prevent the display of common components  
7 (and associated links) that are part of paths with low path strengths. In alternate  
8 embodiments, instead of displaying the relationships between the two end points,  
9 information regarding those relationships is used in another process, such as  
10 filtering search results or generating a report.

11 Fig. 3 is a flow diagram illustrating an embodiment of a procedure 300 for  
12 displaying one or more relationships associated with points (e.g., end points) in an  
13 environment, such as a social environment. Initially, the procedure displays the  
14 first end point (block 302) as well as components associated with the first end  
15 point (block 304). Procedure 300 continues by displaying the second end point  
16 (block 306) and displaying components associated with the second end point  
17 (block 308). These components were previously identified, for example, in the  
18 procedure of Fig. 2 discussed above.

19 At block 310, at least one common component between the first and second  
20 end points is displayed. Additionally, some or all of the links or paths associated  
21 with the common component are displayed (e.g., the link between the first end  
22 point and the component, and the link between the second end point and the  
23 component). The number of common components displayed may vary depending  
24 on the total number of common components identified and any configuration  
25 information associated with the system displaying the components. For example,

1 an administrator or other user may limit the display of common components to the  
2 one or two common components having the strongest path between the two end  
3 points. Additionally, some of the links or paths associated with the common  
4 component may not be displayed. For example, links associated with the common  
5 component that are not associated with one of the end points may not be  
6 displayed. By reducing the number of displayed links, the display is simplified  
7 and focuses attention on links that are related to one or both of the end points.

8 Procedure 300 continues by displaying additional information regarding an  
9 end point when a user identifies the end point (block 312). A user can identify an  
10 end point by moving a cursor or pointer over the end point (e.g., “mousing over”  
11 the end point), or otherwise selecting a particular end point in the display. For  
12 example, when a user moves a cursor over a particular end point, a window, pane,  
13 tool tip, or other display element is added to the display (e.g., overlaid on the  
14 existing graphical display) and contains additional information about the particular  
15 end point. This additional information is typically displayed temporarily for a  
16 predetermined time or until the user moves the cursor or pointer away from the  
17 identified end point. For example, if the identified end point is a person, the  
18 additional information displayed regarding the person may include the person’s  
19 full name, a department in which the person works, groups to which the person  
20 belongs, the person’s telephone number and email address, and the location of the  
21 person’s office. In another embodiment, if the identified end point is a group, the  
22 additional information displayed regarding the group may include members of the  
23 group, a primary contact person for the group, a schedule of upcoming group  
24 activities, and other groups associated with the particular group. Alternatively, any  
25

1 type of information regarding an end point may be displayed when a user  
2 identifies the end point.

3         At block 314, additional information is displayed regarding a link when a  
4 user identifies the link. A user can identify a link by moving a cursor or pointer  
5 over the link, or otherwise selecting a link shown in the display. This additional  
6 information is typically displayed temporarily for a predetermined time or until the  
7 user moves the cursor or pointer away from the identified link. For example, if the  
8 identified link is a link between two people, the additional information displayed  
9 regarding the link may include information regarding the two people associated  
10 with the link, the relationship between the two people (e.g., person A manages  
11 person B), and common groups with which both people are associated.

12         Alternatively, any type of information regarding a link may be displayed when a  
13 user identifies the link.

14         At block 316, additional information is displayed regarding any other  
15 component identified by a user. As discussed above, this additional information  
16 may be displayed temporarily in a window, pane, tool tip, or other display  
17 element. For example, a user may identify a person in the display that is not an  
18 end point, a group or department in the display that is not an end point, or any  
19 other component contained in the graphical display. If the identified component is  
20 a person, additional information includes, for example, information regarding the  
21 person such as name, email address, groups with which the person is associated,  
22 and the like. In another example, if the component is a group, the additional  
23 information can include members of the group, other related groups, a manager of  
24 the group, and the like. In other embodiments, any type of information can be  
25 displayed in response to a user identifying other types of components.

1           Fig. 4 illustrates an example graphical display of information 400 related to  
2 multiple points in a social environment. The example of Fig. 4 includes two  
3 endpoints – Point A, identified by reference number 402, and Point B, identified  
4 by reference number 404. In this example, a user indicated a desire to learn of  
5 relationships between Point A and Point B. In one embodiment, the graphical  
6 display of information 400 is generated using the procedure discussed above with  
7 respect to Fig. 3.

8           Point A has an associated social context identified by broken line 406. This  
9 social context identifies other components that are associated with Point A. Point  
10 A social context 406 includes six components 414 coupled to each other and  
11 coupled to Point A as shown by the interconnecting links between the components  
12 and Point A. These interconnecting links represent relationships between  
13 components. As mentioned above, these relationships can be identified using  
14 organizational charts, mailing lists, and the like. As shown in Fig. 4, certain  
15 components 414 in social context 406 are directly connected to one another or  
16 directly connected to Point A, while other components are coupled to other  
17 components or to Point A via one or more intermediate components.

18           Point B has an associated social context identified by broken line 410. This  
19 social context identifies other components that are associated with Point B. Point  
20 B social context 410 includes five components coupled to each other and coupled  
21 to Point B as shown by the interconnecting links between the components and  
22 Point B. As discussed above, these interconnecting links represent relationships  
23 between components and may be identified using organizational charts, mailing  
24 lists, and the like. Similar to social context 406, certain components in social  
25 context 410 are directly connected to one another or directly connected to Point B,

1 while other components are coupled to other components or to Point B via one or  
2 more intermediate components.

3 A third social context is identified by broken line 408. Social context 408  
4 may also be referred to as an overlapping social context. Social context 408  
5 includes a component 412 that is common to both Point A and Point B. Social  
6 context 408 also includes seven other components coupled to one another by the  
7 interconnecting links between the components. As discussed above, these  
8 interconnecting links represent relationships between components. Similar to  
9 social contexts 406 and 410, certain components in social context 408 are directly  
10 connected to one another, while other components are coupled to other  
11 components via one or more intermediate components.

12 As shown in Fig. 4, a strong relationship between Point A and Point B  
13 exists in component 412. This relationship is defined by a link 416 between Point  
14 A and component 412 as well as a link 418 between component 412 and Point B.  
15 In this example, component 412 is the strongest relationship between Point A and  
16 Point B because this single component 412 is related to both Point A and Point B.  
17 Although other components have links to Point A or to Point B, no other  
18 component has a direct link to Point A and Point B. For example, a link 422  
19 connects component 420 with a component in social context 406. However,  
20 component 420 does not have a direct link to Point A or Point B. Therefore,  
21 component 420 is not part of the strongest relationship between Point A and Point  
22 B. Similarly, a link 426 connects component 424 with a component in social  
23 context 408. However, component 424 does not have a direct link to Point A or  
24 Point B. Therefore, component 424 is not part of the strongest relationship  
25 between Point A and Point B.

1 As shown in Fig. 4, each social context is labeled to allow viewers to  
2 understand the common feature(s) of the social context (i.e., how the components  
3 in that social context are related). For example, social context 406 is labeled  
4 “Group 1”, indicating that the components in social context 406 are related to one  
5 another due to their common group membership (i.e., Group 1). Similarly, social  
6 context 408 is labeled “Group 2” and social context 410 is labeled “Group 3”. Use  
7 of such labels is optional. In alternate embodiments, additional information about  
8 the associated group (or other common feature) is displayed in the vicinity of the  
9 social context. If components in a particular social context have multiple common  
10 features, information regarding all of the common features is displayed. In  
11 alternate embodiments, only information regarding the strongest common feature  
12 is displayed. For example, the strongest common feature may be the smallest  
13 mailing list or group list that contains most of the components in the social  
14 context.

15 In this example, if component 412 is a person, Point A and Point B are  
16 related to one another via this person. Alternatively, if component 412 is a group  
17 or a department, Point A and Point B are related to one another via this group or  
18 department. Additionally, other components in social context 408 are shown in  
19 Fig. 4, thereby providing information regarding the relationship of component 412  
20 with other components in social context 408.

21 The appearance of links and components in Fig. 4 can vary to identify  
22 certain characteristics or other information associated with the links and  
23 components. For example, links associated with a strongest relationship between  
24 two end points may be identified using a different color, different type of line (e.g.,  
25 broken instead of solid), or a different line weight or line thickness. This

1 modification of the appearance of a link is generally referred to as “highlighting”  
2 the link. In the example of Fig. 4, links 416 and 418 are shown with a thicker line  
3 than the other links, thereby signifying that those links are associated with the  
4 strongest relationship between Point A and Point B. In other embodiments,  
5 different line colors are used to identify different types of relationships between  
6 components.

7 Fig. 4 represents one example arrangement of components and links  
8 between two endpoints. Alternate embodiments may include additional  
9 components or links, may delete one or more components or links, or may include  
10 additional types of information or other data. For example, names (e.g., a person’s  
11 name or a group’s name) may be displayed adjacent one or more components in  
12 the graphical display. Further, any number of social contexts may be displayed  
13 along with any number of end points. In a particular embodiment, social contexts  
14 406, 408 and 410 each represent a group, project, or department.

15 In one embodiment, the arrangement of components and links in a  
16 graphical display is influenced by the strength of the relationship between various  
17 components. For example, if two components are closely related, the link between  
18 the two components may be shorter or may be highlighted to indicate a strong  
19 link. Similarly, if two components are not closely related (e.g., the relationship is  
20 based on membership in a common group that has many members), the link  
21 between the two components may be longer to indicate a weaker link.

22 Fig. 5 illustrates the example graphical display of Fig. 4 with additional  
23 information displayed regarding a particular point in the social environment. The  
24 additional information is related to Point B and is shown in box 502. This  
25 additional information is shown when a user identifies Point B (e.g., by mousing



1 over the Point B component in the display). In the example of Fig. 5, the  
2 additional information in box 502 is displayed temporarily (e.g., for a particular  
3 period of time or for as long as the user continues to identify Point B). In this  
4 example, Point B is a person and box 502 contains information about that person,  
5 such as their job title, organizational grouping and manager. Box 502 also lists  
6 several groups that the person is associated with, such as a planning group, an  
7 audio/video group and a marketing group. The numbers in parenthesis after each  
8 group name represent the number of people in that particular group. In this  
9 example, the groups are sorted by the number of people in the group in increasing  
10 order. Alternatively, information contained in box 502 may be arranged in any  
11 order. Further, the amount of information displayed in box 502 may be limited or  
12 truncated to fit within a predetermined size associated with box 502.

13 In other embodiments, other group-related data can be displayed such as a  
14 group membership roster, an owner of the group, permissions required to access  
15 the group, a most senior person in the group, how often the group is accessed by  
16 users, the age of the group, the type of group (e.g., social or business), and the  
17 primary constituents of the group (e.g., developers, managers, or salespeople). In  
18 alternate embodiments, when a user identifies a different component or a link  
19 between components, a box similar to box 502 is displayed that contains  
20 additional information related to the identified component or link. For example, if  
21 a link is identified, the box may contain information regarding groups or mailing  
22 lists that are common to the components at opposite ends of the link.

23 Fig. 6 illustrates another example graphical display 600 of information  
24 related to multiple points in a social environment. In this example, components  
25 that are not on a path that couples Point A to Point B are not displayed. Thus,

1 components that are not coupled (either directly or indirectly) to Point A and to  
2 Point B are not displayed. Graphical display 600 provides a simplified display of  
3 information (as compared to the example of Fig. 4) by removing components that  
4 are not contained in a path that relates Point A to Point B.

5 Fig. 7 illustrates another example graphical display 700 of information  
6 related to two people (John and Mark) in a social environment. In this example,  
7 additional information is displayed regarding the management chain associated  
8 with the two people. John is associated with his manager 702, illustrated by a link  
9 704. John's manager 702 is associated with their manager 708, illustrated by a  
10 link 708. Similarly, Mark is associated with his manager 710, illustrated by a link  
11 712. Mark's manager 710 is associated with manager 706, illustrated by a link  
12 714. Thus, Fig. 7 illustrates management relationships as well as other social  
13 relationships between two people. In alternate embodiments, other information,  
14 relationships, and the like are displayed regarding two or more components in a  
15 social environment.

16 Particular examples discussed herein illustrate paths between two end  
17 points having two links – one link from a first end point to a common component  
18 and a second link from the common component to a second end point. In many  
19 situations, the shortest path (i.e., the fewest links and/or the fewest intermediate  
20 components) is the most meaningful relationship. However, in alternate  
21 embodiments, paths may contain any number of links (and any number of  
22 intermediate components) between two end points. For example, other criteria  
23 (such as persons in a particular group or persons that are managers) may be more  
24 important than the number of intermediate components.  
25

1        Fig. 8 illustrates a general computer environment 800, which can be used to  
2        implement the techniques described herein. The computer environment 800 is  
3        only one example of a computing environment and is not intended to suggest any  
4        limitation as to the scope of use or functionality of the computer and network  
5        architectures. Neither should the computer environment 800 be interpreted as  
6        having any dependency or requirement relating to any one or combination of  
7        components illustrated in the example computer environment 800.

8        Computer environment 800 includes a general-purpose computing device in  
9        the form of a computer 802. The components of computer 802 can include, but  
10       are not limited to, one or more processors or processing units 804 (optionally  
11       including a cryptographic processor or co-processor), a system memory 806, and a  
12       system bus 808 that couples various system components including the processor  
13       804 to the system memory 806.

14       The system bus 808 represents one or more of any of several types of bus  
15       structures, including a memory bus or memory controller, a point-to-point  
16       connection, a switching fabric, a peripheral bus, an accelerated graphics port, and  
17       a processor or local bus using any of a variety of bus architectures. By way of  
18       example, such architectures can include an Industry Standard Architecture (ISA)  
19       bus, a Micro Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a  
20       Video Electronics Standards Association (VESA) local bus, and a Peripheral  
21       Component Interconnects (PCI) bus also known as a Mezzanine bus.

22       Computer 802 typically includes a variety of computer readable media.  
23       Such media can be any available media that is accessible by computer 802 and  
24       includes both volatile and non-volatile media, removable and non-removable  
25       media.

1       The system memory 806 includes computer readable media in the form of  
2 volatile memory, such as random access memory (RAM) 810, and/or non-volatile  
3 memory, such as read only memory (ROM) 812. A basic input/output system  
4 (BIOS) 814, containing the basic routines that help to transfer information  
5 between elements within computer 802, such as during start-up, is stored in ROM  
6 812. RAM 810 typically contains data and/or program modules that are  
7 immediately accessible to and/or presently operated on by the processing unit 804.

8       Computer 802 may also include other removable/non-removable,  
9 volatile/non-volatile computer storage media. By way of example, Fig. 8  
10 illustrates a hard disk drive 816 for reading from and writing to a non-removable,  
11 non-volatile magnetic media (not shown), a magnetic disk drive 818 for reading  
12 from and writing to a removable, non-volatile magnetic disk 820 (e.g., a “floppy  
13 disk”), and an optical disk drive 822 for reading from and/or writing to a  
14 removable, non-volatile optical disk 824 such as a CD-ROM, DVD-ROM, or other  
15 optical media. The hard disk drive 816, magnetic disk drive 818, and optical disk  
16 drive 822 are each connected to the system bus 808 by one or more data media  
17 interfaces 825. Alternatively, the hard disk drive 816, magnetic disk drive 818,  
18 and optical disk drive 822 can be connected to the system bus 808 by one or more  
19 interfaces (not shown).

20       The disk drives and their associated computer-readable media provide non-  
21 volatile storage of computer readable instructions, data structures, program  
22 modules, and other data for computer 802. Although the example illustrates a hard  
23 disk 816, a removable magnetic disk 820, and a removable optical disk 824, it is to  
24 be appreciated that other types of computer readable media which can store data  
25 that is accessible by a computer, such as magnetic cassettes or other magnetic

1 storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or  
2 other optical storage, random access memories (RAM), read only memories  
3 (ROM), electrically erasable programmable read-only memory (EEPROM), and  
4 the like, can also be utilized to implement the example computing system and  
5 environment.

6 Any number of program modules can be stored on the hard disk 816,  
7 magnetic disk 820, optical disk 824, ROM 812, and/or RAM 810, including by  
8 way of example, an operating system 826, one or more application programs 828,  
9 other program modules 830, and program data 832. Each of such operating  
10 system 826, one or more application programs 828, other program modules 830,  
11 and program data 832 (or some combination thereof) may implement all or part of  
12 the resident components that support the distributed file system.

13 A user can enter commands and information into computer 802 via input  
14 devices such as a keyboard 834 and a pointing device 836 (e.g., a "mouse").  
15 Other input devices 838 (not shown specifically) may include a microphone,  
16 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and  
17 other input devices are connected to the processing unit 804 via input/output  
18 interfaces 840 that are coupled to the system bus 808, but may be connected by  
19 other interface and bus structures, such as a parallel port, game port, or a universal  
20 serial bus (USB).

21 A monitor 842 or other type of display device can also be connected to the  
22 system bus 808 via an interface, such as a video adapter 844. In addition to the  
23 monitor 842, other output peripheral devices can include components such as  
24 speakers (not shown) and a printer 846 which can be connected to computer 802  
25 via the input/output interfaces 840.

1 Computer 802 can operate in a networked environment using logical  
2 connections to one or more remote computers, such as a remote computing device  
3 848. By way of example, the remote computing device 848 can be a personal  
4 computer, portable computer, a server, a router, a network computer, a peer device  
5 or other common network node, game console, and the like. The remote  
6 computing device 848 is illustrated as a portable computer that can include many  
7 or all of the elements and features described herein relative to computer 802.

8 Logical connections between computer 802 and the remote computer 848  
9 are depicted as a local area network (LAN) 850 and a general wide area network  
10 (WAN) 852. Such networking environments are commonplace in offices,  
11 enterprise-wide computer networks, intranets, and the Internet.

12 When implemented in a LAN networking environment, the computer 802 is  
13 connected to a local network 850 via a network interface or adapter 854. When  
14 implemented in a WAN networking environment, the computer 802 typically  
15 includes a modem 856 or other means for establishing communications over the  
16 wide network 852. The modem 856, which can be internal or external to computer  
17 802, can be connected to the system bus 808 via the input/output interfaces 840 or  
18 other appropriate mechanisms. It is to be appreciated that the illustrated network  
19 connections are exemplary and that other means of establishing communication  
20 link(s) between the computers 802 and 848 can be employed.

21 In a networked environment, such as that illustrated with computing  
22 environment 800, program modules depicted relative to the computer 802, or  
23 portions thereof, may be stored in a remote memory storage device. By way of  
24 example, remote application programs 858 reside on a memory device of remote  
25 computer 848. For purposes of illustration, application programs and other

1 executable program components such as the operating system are illustrated herein  
2 as discrete blocks, although it is recognized that such programs and components  
3 reside at various times in different storage components of the computing device  
4 802, and are executed by the data processor(s) of the computer.

5 Various modules and techniques may be described herein in the general  
6 context of computer-executable instructions, such as program modules, executed  
7 by one or more computers or other devices. Generally, program modules include  
8 routines, programs, objects, components, data structures, etc. that perform  
9 particular tasks or implement particular abstract data types. Typically, the  
10 functionality of the program modules may be combined or distributed as desired in  
11 various embodiments.

12 An implementation of these modules and techniques may be stored on or  
13 transmitted across some form of computer readable media. Computer readable  
14 media can be any available media that can be accessed by a computer. By way of  
15 example, and not limitation, computer readable media may comprise “computer  
16 storage media” and “communications media.”

17 “Computer storage media” includes volatile and non-volatile, removable  
18 and non-removable media implemented in any method or technology for storage  
19 of information such as computer readable instructions, data structures, program  
20 modules, or other data. Computer storage media includes, but is not limited to,  
21 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,  
22 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic  
23 tape, magnetic disk storage or other magnetic storage devices, or any other  
24 medium which can be used to store the desired information and which can be  
25 accessed by a computer.

1           “Communication media” typically embodies computer readable  
2 instructions, data structures, program modules, or other data in a modulated data  
3 signal, such as carrier wave or other transport mechanism. Communication media  
4 also includes any information delivery media. The term “modulated data signal”  
5 means a signal that has one or more of its characteristics set or changed in such a  
6 manner as to encode information in the signal. By way of example, and not  
7 limitation, communication media includes wired media such as a wired network or  
8 direct-wired connection, and wireless media such as acoustic, RF, infrared, and  
9 other wireless media. Combinations of any of the above are also included within  
10 the scope of computer readable media.

11           Although the description above uses language that is specific to structural  
12 features and/or methodological acts, it is to be understood that the invention  
13 defined in the appended claims is not limited to the specific features or acts  
14 described. Rather, the specific features and acts are disclosed as exemplary forms  
15 of implementing the invention.